

Multicast Data Transfer

Field of the Invention

5 The invention relates to the simultaneous delivery of information to multiple hosts. In particular, the invention relates to an apparatus and method for transferring text, audio, video or other data to multiple hosts in a communication system.

Background

Methods for transmitting data from a source to multiple users over a network fall 10 into one of three categories. Unicast is a point-to-point delivery mechanism, where a source transmits a separate stream of data for each user. The volume of data transmitted increases linearly with the number of users, leading to the consumption of an unacceptably high level of resources where the number of recipients is large. The remaining two categories define point-to-multipoint or multipoint-to- 15 multipoint delivery mechanisms. Broadcasting refers to the transmission of data to all hosts on a network. While this may work well in small networks, it requires the replication of a large number of data packets when used in a wide area network with many users. Furthermore, this method can be wasteful where the data is not of interest to a significant proportion of the recipients.

20 The third category is multicasting, which is the simultaneous transmission of data to a select group of users. Multicasting has particular application in fields requiring streaming of content to multiple users in real time, such as news feeds, online gaming, Digital Video Broadcasting (DVB) and videoconferencing. Overviews of this technique are given in "Multicast Networking and Application" by C. Kenneth 25 Miller, Addison-Wesley 1988 [ISBN 0-201-30979-3] and in "Deploying IP MULTICAST in the Enterprise" by T. Maufer, Prentice Hall PTR, 1998 [ISBN 0-13-897687-2]. The use of multicasting in a cellular radio access network is discussed in "Multimedia Broadcast/ Multicast Service", 3rd Generation Partnership Project Technical Specification 3GPP™ TS 22.146 v1.0.0 2001.

Figure 1 is a block diagram of a prior communication system 1, for multicasting information from a content provider, i.e., host 2, to a plurality of receiver hosts 3 via the Internet 4. For simplicity, the data path 5 is depicted as extending directly 5 from the content provider 2 to a router 6 without intervening stages, such as other transmitters or routers, which may be located between these nodes. The data paths leading from the content provider 2 to the receiver hosts 3a-3c diverge at the router 6. Situations where this divergence include those where the receiver hosts 3 do not 10 belong to the same sub-network, where the receiver hosts 3 are served by different transmitters 7 or where the radio access network is not multicast enabled.

Instead of sending a separate stream of data packets A-D for each individual receiver host 3a, 3b, the content provider 2 provides a single data stream 5. The data stream is replicated by a multicast-enabled edge router 6 wherever the paths to 15 the different receiver hosts 3a, 3b diverge. For example, in Figure 1, the path from the content provider 2 to receiver hosts 3a, 3b is common to both hosts. The router 6 copies the incoming data packets A-E, producing duplicate data streams 8a, 8b for transmission to the receiver hosts 3a, 3b. As the content provider 2 transmits only a single data stream 5, the use of its system resources and the 20 network load are reduced when compared with a unicast system. Furthermore, as the transmission of replicated data is kept to a minimum, multicasting may be used in wide area networks where broadcasting is unfeasible.

There are two types of multicasting services. In the first, a content provider has a 25 reserved bandwidth for delivering a predetermined service to a group of end-users. The second type allows end-users to select the service provided to them. In both cases, a user can receive the multicast stream of data by submitting a request to join the group by sending a message to the content provider in a format specified by the Internet Engineering Task Force (IETF). The user can leave the multicast group by 30 submitting a corresponding request.

Multicasting is a suitable mechanism for providing end-users with a continuous service and has limited applicability for "one-shot" transmissions. For example, a

group of users may require the delivery of a file, for example, a multimedia clip, a web page, an mp3 file, a document or a software module for an application or a game. Referring again to Figure 1, the users of receiver hosts 3a, 3b may both request the file, which comprises data packets A-E, from the content provider 2. As 5 the data stream 5 is received by the router 6, the data packets are copied and forwarded to a radio transmitter 7, which transmits the encapsulated data packets to receiver hosts 3a, 3b as data streams 8a, 8b over a wireless network.

There are currently no provisions for allowing a host to join an ongoing file delivery 10 transmission, so a third host 3c, requesting a file after the start of the transmission, would be excluded. Even if the third host 3c were to join the file delivery transmission, data packets A-C would have been missed and it is probable that the user of host 3c would require the entire file. This exclusion increases the likelihood that a transmission is received by a low number of recipients, or even a single user. 15 Any requests submitted after the delivery of the data packets A-E has commenced are accommodated by repeating the file delivery transmission, reducing the efficiency gain associated with the use of multicasting and increasing the use of the air interface bandwidth. As the use of the air interface in providing Internet type services rises, the efficient use of air interface bandwidth will become increasingly 20 important.

Prior systems have addressed these problems by defining a schedule of forthcoming transmissions. A user can subscribe to a delivery list for a particular file, or set of files, and receive the relevant data at the next scheduled opportunity. However, this 25 delivery mechanism does not respond directly to the user's requirements and is inflexible as, again, once the file delivery transmission has begun, any other further users requiring the same file are excluded and must subscribe to the next scheduled transmission.

30 Summary of the Invention

An object of the invention is to allow multicast file delivery in a manner that is more responsive to users' needs, with particular application to scalable multicasts,

such as unidirectional IP multicasts. The invention has particular advantages in systems comprising a wireless communications network by reducing use of the air interface bandwidth.

- 5 In a first aspect of the invention, a multicast-enabled network element comprises a first logical interface for receiving data from first host, a second logical interface for transmitting said data to one or more further hosts, a processor for defining a group comprising one or more further hosts, wherein a further host is added to the group in response to the reception of a request and a cache, wherein the network element
10 is configured to store received data in the cache until a predetermined condition is met and, in response to the meeting of this condition, to forward the data to said further hosts in said group and the processor is configured to limit the group to further hosts situated at the same location.
- 15 In a second aspect of the invention, a method of file delivery over a network comprises the steps of receiving a request for the file from a first host at a network element, retrieving the file from a second host, storing the file in a cache associated with the network element, defining a group including the first host, waiting for a period of time until a predetermined condition is met where, if further requests for
20 said file are received by the network element from one or more other hosts before the period of time expires, said one or more other hosts are added to the group, and forwarding the file to the first host and to any other hosts in said group, wherein the group is limited to the first host and other hosts situated at the same location as the first host.
- 25 In a third aspect of the invention, a multicast-enabled network element comprises a cache, said network element being configured to perform the following steps in response to a request for the delivery of a file from a first host: retrieve the requested file from a second host, store said file in said cache, define a group
30 including the first host, delay forwarding the file to the first host for a period of time until a predetermined condition is met and, if further requests for said file are received by the network element from one or more other hosts before the period of time expires, add said one or more other hosts to the group and forward the file to

the first host and to any other hosts in said group and is further configured to limit the group to the first host and to other hosts situated at the same location as the first host.

- 5 In a fourth aspect of the invention, a network element comprises first receiving means for receiving data from a first host, second receiving means for receiving a request from one or more further hosts, means for defining a group of further hosts, wherein a further host is added to the group in response to the reception of a request, forwarding means for forwarding said data to the further hosts and data storage means, wherein the network element is configured to store received data in the data storage means until a predetermined condition is met and, in response to the meeting of this condition, to retrieve said data and to forward the data to the further hosts in the group.
- 10
- 15 The implementation of the invention in a network element at the air interface, such as an edge router and the compilation of a group of file receiving hosts defined in terms of their location, and in particular the communications cell in which they are located, results in a reduction in use of the air interface bandwidth.
- 20 The use of a predetermined condition relating to the number of requests is intended to increase the likelihood that a given transmission will be received by more than one end user. The second condition assigns a maximum time period for the receipt of further requests. This places a maximum limit on the waiting time that may be experienced by a user, so that, for example, a single user is not subjected to an indefinite delay before receiving the requested file. The maximum time period may be fixed or may change dynamically in response to the requirements of the user and/ or the network.
- 25
- 30 Where a host submits a file request during this time period, the time interval between the submission of the request and the receipt of the complete file is considerably reduced, increasing the data transmission rate perceived by such a user and thereby improving the quality of service.

Brief Description of the Drawings

An embodiment of the invention will now be described with reference to the accompanying drawings, in which:

5

Figure 1 is a block diagram of a prior art communication system;

Figure 2 is a block diagram of a communication system in accordance with an embodiment of the present invention;

10 Figure 3 is a flowchart of a method according to an embodiment of the present invention as enacted by a router;

Figure 4 is a block diagram of a router in accordance with an embodiment of the present invention; and

15 Figure 5 depicts a mobile handset terminal for use according to an embodiment of the invention.

15

Detailed Description

Figure 2 depicts an embodiment of a communication system 9 for multicast file delivery via the Internet 4 comprising a multimedia content provider 2, receiver hosts 3a-3c, a router 6, a radio transmitter 7 and a cache 10. The data path 5 may include further elements, such as other transmitters or routers, which may be located between the content provider 2 and router 6. The portion of the data paths leading from the content provider 2 to the receiver hosts 3a-3c have coinciding portions 5 extending between the content provider 2 and the router 6 and diverge 20 between the router 6 and the individual receiver hosts 3a-3c.

25 In this particular example, receiver hosts 3a-3c are mobile devices belonging to a cellular wireless DVB-T wide area network (WAN) and the router 6 is an edge router, i.e. the last router before the air-interface in the file delivery path between the content provider 2 and the receiver hosts 3a-3c. The router 6 forwards the 30 multicast data to the radio transmitter 7, which sends the data to the receiver hosts over the wireless network in accordance with a suitable communication protocol. Suitable protocols include, but are not limited to, the following protocols: Reliable

Multicast Transport Protocol (RMTP), Reliable Multicast File Transfer Protocol (RMFTP), Asynchronous Layered Coding (ALC), NACK-Oriented Reliable Multicast (NORM), Pragmatic General Multicast (PGM), Tree Acknowledgement based protocol (TRACK), User Datagram Protocol (UDP) and Unidirectional Hypertext Transfer Protocol (UHTTP). The Session Announcement Protocol (SAP) may also be used for sending service information to the receiver hosts 3a-3c.

An embodiment of the file delivery procedure followed by the router 6 is described with reference to Figure 3. Beginning at step s0, the router 6 receives a request for a file from one of the receiver hosts, 3a (step s1), which is situated at a first location. The request includes the address information for the receiver host 3a, which is extracted by the router 6 for use in addressing data packets and messages directed to that receiver host 3a.

15 The file may be, for example, a multimedia clip of a goal in a football or hockey match. A number of end users, aware that a goal has been scored through other information sources, such as a radio broadcast, may submit a request for a clip within a short time interval following a goal.

20 In this example, where the multicast enabled network is unidirectional, communications between the receiver host 3a and the router 6 are made via another network or communication system 11, such as the Internet or a cellular telecommunications network. The router 6 requests the file from the content provider 2 (step s2), which responds by transmitting the file and an associated delivery message to the router 6. The delivery message includes information about the time the file was created and an indication of when it will next be updated.

25 In a conventional system, the router 6 would forward the data packets A-G to the receiver host 3a immediately and, once the transmission has commenced, it would not be possible for any further receiver hosts 3b, 3c, to join the multicast group and receive the data packets in that file delivery transmission.

To overcome this problem, the present router 6 receives the incoming data from the content provider 2 (step s3) and stores it in the cache 10 (step s4). The router 6 is configured to check that the file stored in the cache 10 is the latest version available (step s5) using the delivery message.

5

A file delivery list is defined comprising an entry corresponding to the first receiver host 3a (step s6) and a timer t is initialised (step s7). The delivery of the file is then delayed by a time period T. Notification that the delivery of the file is expected to commence after the expiry of time period T is sent to the first receiver host 3a. The 10 time period T, which is divided into time intervals t1 (step s8), is allowed for the reception of any further requests for the same file from other receiver hosts, e.g., 3b, 3c at the same location as the first receiver host 3a, so that the receiver hosts 3a- 15 3c can receive the file in the same file delivery transmission. The time period T can be fixed or may change dynamically, for example, so that a shorter period may be allowed as the number of received requests increases. Alternatively, the period T 20 may depend on the time of the next update, as indicated in the delivery message, so that an updated file is transmitted to the receiver hosts 3, or on the size of the data file, so that a longer period is allowed for larger files in order to increase the likelihood of accommodating the requests of a number of end-users with a single file delivery transmission.

25

If a further request is received (step s9), the router 6 determines whether the receiver host 3b submitting the request is situated at the same location as the first receiver host 3a (step s10). The router 6 considers the receiver hosts 3a, 3b to be at the same location if they are within an area covered by the same cell, the cell being defined in either a bi-directional telecommunications network, where such a network is used for submitting requests to the router 6, or the wireless DVB-T network used for file delivery. Where the host has access to data about its location, such as its DVB-T cell, it will include this information in the request for the file for this purpose. If the receiver host 3b is at the same location, it is added to the file delivery list (step s11). If not, the router 6 then adds an entry to a separate file delivery list (step s12) corresponding to the location of receiver host 3b, setting up a 30

new file delivery list if necessary. The receiver host 3b is then notified of the expected start time of the file delivery transmission, i.e. the end of time period T.

Any further requests submitted by receiver hosts 3b, 3c at the first location are 5 handled by the router 6 and should not be directed to the content provider 2. However, the content provider 2 has the facility to check incoming requests as in case such a request is received. The content provider 2 compares incoming requests and, if a request from a host 3b matches a previous request from another host 3a in terms of the file requested and the location of the receiver host 3, information 10 relating to the match is sent to the router 6, so that the host 3b that submitting the later request can be added to the file delivery list.

After a further request has been received, or after the expiry of time interval t1, the router 6 checks whether the number of requests entered in the file delivery list 15 exceeds a predetermined threshold N (step s13). If this condition is not met and the period T has not yet elapsed (step s14), the router waits for a further interval t1 (step s8). This process is repeated until one of these conditions is satisfied, i.e., either the number of requests exceeds the threshold or the time period T has elapsed.

20 The router 6 then retrieves the data packets A-G from the cache 10 (step s15). Where more than one end-user has requested the file, the router 6 copies the data packets A-G, producing multiple data streams 8a-8c and transmits them to the receiver hosts 3a-3c (step s16). The file delivery procedure is then complete (step 25 s17).

An example of a suitable router 6 is shown in Figure 4. In addition to the cache 10, the router 6 comprises an input/output interface 12 for receiving data from the content provider 2, input/output interfaces 13-15 for communicating with the 30 receiver hosts 3a-3c and a number of other components connected by a data bus 16 as follows. A processor 17 monitors and controls routing operations, including generating processes for establishing and releasing connections between the router 6 and receiver hosts 3a-3c. A memory facility 18 is provided for storing the router

application software. An address processor 19 extracts the address information from incoming data packets and uses this information to determine how each data packet is to be processed, e.g., whether an incoming packet is a request for a file from a receiver host 3a-3c or file data. Incoming file data is stored in the cache 10, although the router 6 may also use an external buffer for this purpose. Requests are forwarded to a file request handler 20, which stores the address of the receiver host 3 making the request in a file delivery list and manages the copying of data packets A-G for transmission to the receiver hosts 3a-3c. The file request handler 20 also controls any queuing processes, e.g. where a further request has been received after the file delivery transmission has commenced. A clock 21 is associated with the processor 17 for controlling the timer functions described above.

The receiver hosts 3 may be mobile telephone handset terminals 22, as shown in Figure 5. The mobile telephone handset terminal comprises an antenna 23, a transceiver 24, a central processing unit (CPU) 25 and a battery 26. The transceiver 24, which is, for example, a GSM, GPRS, 3G, or similar bi-directional transceiver, is connected to the antenna 23 and the CPU 25. Also connected to the CPU 25 are a receiver (DVB-T) 27, a keypad 28 and a display 29. The terminal 22 also includes other conventional features of mobile telephone handsets, but these are omitted from Figure 5 for the sake of clarity.

Where three receiver hosts 3a-3c receive the same file in a single multicast, the increase in efficiency over a unicast system, or the prior multicast system where the likelihood of a transmission being received by a single receiver host 3 is significantly increased, is three-fold, i.e., the data packets A-G are sent through the network 4 once, instead of three times. However, as the transmission to the first receiver host 3a has been delayed by time T, or by the time taken to receive N requests, the user of host 3a will perceive an associated decrease in bandwidth, but this may be offset by "pooling" of the bandwidth available to the three receiver hosts 3a-3c

The performance of the present system is compared with a prior art system in the following example.

In the prior art system, the file is delivered to the first host 3a immediately and, as the subsequent requests from hosts 3b, 3c cannot be met from the ongoing transmission, the other hosts 3b, 3c receive the data in a subsequent file delivery transmission. For example, downloading a file of 10 Mbits over a network with an available bandwidth of 1 Mbit/s would take 10 s. A file delivery transmission sending data to receiver host 3a begins at 0 s. Requests from the users of hosts 3b and 3c are submitted after the file delivery transmission to host 3a has commenced, at 4 s and 12 s respectively. The resulting waiting times and perceived bandwidths in the prior system would be 10 s and 1 Mbit/s for host 3a, 14 s and 700 kbit/s for host 3b. The request from host 3c would be received after the second file delivery transmission, to host 3b, has started, and so would be queued until a third file delivery transmission could begin. Host 3c would therefore experience a waiting time of 18 s and perceive a bandwidth of 550 kbit/s. Therefore, under the prior system, the three hosts 3a-3c would perceive an average bandwidth of 750 kbit/s, if the method of the present invention is followed, these time periods and bandwidths are 25 s and 400 kbit/s for host 3a, 21 s and 480 kbit/s for host 3b and 13 s and 770 kbit/s for host 3c. This provides an average perceived bandwidth of 550 kbit/s, although this does not reflect the efficiency gains arising from a reduced network load and reduced duplication of transmissions.

The perceived average bandwidth in the system of the present invention further increases with respect to the prior system as more hosts are added to the multicast group. This is demonstrated in Table 1, where t_d indicates the time taken to complete the delivery of the file from the time that User 1 submits their request. PBW denotes the bandwidth perceived by the user.

User	Time of Request (s)	File Size (Mbit)	Transfer Rate (Mbit/s)	Prior System		Present system		
				t_d (s)	PBW (kbit/s)	Waiting period T (s)	t_d (s)	PBW (kbit/s)
1	0	10	1	10	1000	15	25	400
2	4	10	1	20	625	15	21	476
3	12	10	1	30	556	15	13	769
4	12	10	1	40	357	15	13	769
5	12	10	1	50	263	15	13	769
6	13	10	1	60	213	15	12	833
7	13	10	1	70	175	15	12	833
8	14	10	1	80	133	15	11	909
9	15	10	1	90	386	15	10	1000

Table 1

5 The invention has been described by way of an example and can be used for transferring data packets other than IP packets. While a point-to-multipoint transmission has been described, the invention may also be used for multipoint-to-multipoint data transfer. The receiver hosts 3a-3c may be fixed or mobile devices, e.g., mobile telephones or personal digital assistants (PDAs). It is not necessary for 10 the receiver hosts 3a-3c to belong to the same local network or for the location of the hosts to be defined in terms of cell coverage.

Furthermore, it is not necessary for the router 6 to transmit data to the receiver hosts 3a-3c using a wireless communication system. The embodiment described 15 relates to a system comprising a WAN, however, the invention is equally applicable to LAN networks or any other multicast enabled network and may be used for the transfer of files over networks other than the Internet or a DVB-T network, e.g., over DVB-S (satellite), DVB-C (cable), other DVB variants, Integrated Services Digital Broadcasting (ISDB), ATSC Digital Television or Digital Audio 20 Broadcasting (DAB) networks or in multicast-enabled General Packet Radio Service (GPRS), Enhanced Data GSM Environment (EDGE), Universal Mobile

Telecommunications Systems (UMTS) or Wireless Code Division Multiple Access (W-CDMA) or CDMA2000 systems.

Furthermore, although the described embodiment comprised an edge router 6, the 5 invention may be implemented at any network element, e.g., a server, node or router, along the path between a content provider 2 and receiver hosts 3a-3c.